arXiv:1603.00351v1 [stat.CO] 1 Mar 2016

Analyzing Non-proportional Hazards: Use of the MRH Package

Yolanda Hagar and Vanja Dukic *

Abstract

In this manuscript we demonstrate the analysis of right-censored survival outcomes using the MRH package in R. The MRH package implements the multi-resolution hazard (MRH) model ([4, 3, 7, 6, 5]), which is a Polya-tree based, Bayesian semi-parametric method for flexible estimation of the hazard rate and covariate effects. The package allows for covariates to be included under the proportional and non-proportional hazards assumption, and for robust estimation of the hazard rate in periods of sparsely observed failures via a "pruning" tool.

Key Words: MRH, multi-resolution hazard, non-proportional hazards, sparse observations, survival analysis.

1 Introduction

The hazard rate, defined as $h(t) = \lim_{\Delta \to 0} P(t \leq T < t + \Delta \mid T \geq t)/\Delta = f(t)/S(t)$ (where $S(t)$ is the survival function for T and $f(t) = -S'(t)$), can be critical in assessing how the risk of a disease changes over time. However, it can be difficult to estimate reliably, particularly over the course of a study with few observed failures during the follow up period and when the effects of covariates change over time. (For examples, see [2, 27, 21, 16].)

To this end, the MRH package in R ([12, 13]) has three overarching features:

1. Estimation of the hazard rate and the associated credible intervals (as well as the corresponding survival function and cumulative hazard estimates and credible intervals).

2. Joint estimation of the effects of predictors incorporated under the proportional hazards (PH) and non-proportional hazards (NPH) assumptions.

3. A "pruning" tool that combines portions of the hazard rate that are similar, providing robust estimates through periods of sparse failures, and allowing for faster computation times.

*Yolanda Hagar is a postdoctoral researcher in Applied Mathematics, University of Colorado at Boulder. Vanja Dukic is a Professor in Applied Mathematics, University of Colorado at Boulder. Correspondence emails: yolanda.hagar@colorado.edu, vanja.dukic@colorado.edu.

The underlying statistical approach employed in the MRH package is the multi-resolution hazard (MRH) model, a Bayesian semi-parametric hazard rate estimator previously presented and used in [4, 3, 7, 6, 5, 11, 10], and compared to other packages in [13]. This model for survival data is based on the Polya tree methodology, and is flexibly designed for multi-resolution inference capable of accommodating periods of sparse events and varying smoothness. The MRH model accommodates both proportional and non-proportional effects of predictors over time and also uses a pruning algorithm presented in [5] and [11], which performs data-driven "pre-smoothing" of the hazard rate by merging time intervals with similar hazard levels. Pruning has been shown to increase computational efficiency and reduce overall uncertainty in hazard rate estimation in the presence of periods with smooth hazard rate and low event counts ([5]).

The following sections of this manuscript are organized as follows: In Section 2 we provide background on the multi-resolution hazard (MRH) model. In Section 3 we briefly discuss the tongue cancer data set we use to demonstrate the MRH package, and in Section 4 we cover use of the MRH package, including fitting and plotting fitted models, pruning the prior, assessing convergence of the MCMC chains, and the effects adjustment of the prior parameters. Lastly, we discuss our conclusions in Section 5.

2 Multi-resolution hazard modeling

The MRH model is a Bayesian, semi-parametric survival model that produces an estimate of the hazard rate and covariate effects. The MRH prior is closely related to the Polya tree prior ([8, 20]), which is an infinite, recursive, dyadic partitioning of a measurable space Ω. (In practice, this process is terminated at a finite level M.) The MRH prior is a type of Polya tree in that it uses a fixed, pre-specified partition and controls the hazard level within each bin through a multi-resolution parameterization.

To facilitate the recursive dyadic partition of the multiresolution tree, we assume that $J = 2^M$. Here, M is an integer, set to achieve the desired time resolution, or through model selection criteria or clinical input (for example, see [4, 6]). The hazard function is parametrized by a set of hazard increments $d_j, j = 1, \ldots, J$. where d_j represents the aggregated hazard rate over the j^{th} time interval, ranging from (t_{j-1}, t_j). In standard survival analysis notation, $d_j = \int_{t_{j-1}}^{t_j} h(s)ds \equiv H(t_j) - H(t_{j-1})$, where $h(t)$ is the hazard rate at time t. The cumulative hazard, H, is equal to the sum of all 2^M hazard increments, which are denoted as $d_j, j = 1, \ldots 2^M$. The model then recursively splits H at different branches via the "split parameters" $R_{m,p} = H_{m,2p}/H_{m-1,p},\ m = 1, 2, \ldots, M-1,\ p = 0, \ldots, 2^{m-1}-1$. Here, $H_{m,q}$ is recursively defined as $H_{m,q} \equiv H_{m+1,2q} + H_{m+1,2q+1}$ (with $H_{0,0} \equiv H$, and $q = 0, \ldots, 2^m - 1$). The $R_{m,p}$ split parameters, each between 0 and 1, guide the shape of the *a priori* hazard rate over time. The complete hazard rate prior specification is obtained via priors placed on all tree parameters: a Gamma(a, λ) prior is placed on the cumulative hazard H, and Beta prior on each split parameter $R_{m,p}$, $\mathcal{B}e(2\gamma_{m,p}k^m a, 2(1 -$

$\gamma_{m,p})k^m a$). This parametrization ensures the self-consistency of the MRH prior at multiple resolutions ([4, 5]).

The basic MRH model was extended in [7] into the hierarchical multi-resolution (HMRH) hazard model, capable of modeling non-proportional hazard rates in different subgroups jointly with other proportional predictor effects. The pruning methodology for combining similar hazard bins was developed in [5] for individual hazard rates, and combined with the HMRH model in [11]. The pruning algorithm detects consecutive time intervals where failure patterns are statistically similar, increasing estimator efficiency and reducing computing time. The resulting method produces computationally stable and efficient inference, even in periods with sparse numbers of failures, as may be the case in studies with long follow-up periods.

2.1 MRH likelihood function

We denote T_i as the minimum of the observed time to failure or the right-censoring time for subject i. Each subject belongs to one of the $\mathcal{L}$ covariate strata, and within each stratum we employ the proportional hazards assumption such that:

$$h_\ell(t \mid X, \vec{\beta}) = h_{base,\ell}(t) \exp\{\mathbf{X}'\vec{\beta}\}.$$

Here, h_ℓ denotes the baseline hazard rate for treatment strata ℓ, $\mathbf{X}$ represents the $z \times n_\ell$ matrix of z covariates (other than those used for stratification) for the n_ℓ patients in the stratum ℓ, while $\vec{\beta}$ denotes the $z \times 1$ vector of the covariate effects.

For subject i in stratum ℓ with failure time at $T_i \in [0, t_J)$, the likelihood contribution is:

$$L_{i,\ell}(T_i \mid X_i, \vec{\beta}) = h_{base,\ell}(T_i) \exp(X_i'\vec{\beta})^{\delta_i} S_{base,\ell}(T_i)^{\exp(X_i'\vec{\beta})},$$

where X_i is that subject's covariate vector, $S_{base,\ell}$ is the baseline survival function for the stratum ℓ, and δ_i is the censoring indicator that equals 1 if subject i had an observed failure, and 0 otherwise. Thus, the log-likelihood for all n patients in all $\mathcal{L}$ strata together ($n = \sum_{\ell=1}^{\mathcal{L}} n_\ell$) is

$$\log L(\mathbf{T} \mid \tilde{\beta}, \mathbf{H}, \mathbf{R_{m,p}}, \mathbf{X}) = \sum_{\ell=1}^{\mathcal{L}} \sum_{i \in \mathcal{S}_\ell} \left\{ \delta_i \log\left(h_{base,\ell}(T_i) + X_i'\vec{\beta}\right) - \exp\left(X_i'\vec{\beta}\right) H_{base,\ell}(T_i) \right\},$$

where $\mathcal{S}_\ell$ denotes the set of indices for subjects belonging to the stratum ℓ, and $H_{base,\ell}(T) = -\log S_{base,\ell}(T)$. In this model, the $\mathcal{L}$ hazard rates are estimated jointly with all the covariate effects. The non-proportional covariate effect is then calculated as the log of the hazard ratio between different covariate strata in each bin.

2.2 Pruning the MRH model

The MRH prior resolution is often chosen as a compromise between the desire for detail in the hazard rate, and the number of observed (i.e. uncensored) fail-

ures. As the resolution increases, the number of observed failures within each bin decreases. While useful for revealing detailed patterns, a large number of intervals is a large number of model parameters, which will generally require longer computing times and may result in estimators with lower statistical efficiency ([5]).

"Pruning" starts with the full MRH tree prior, and merges adjacent bins that are constructed via the same split parameter, $R_{m,p}$, when the hazard increments in these two bins ($H_{m+1,2p}$ and $H_{m+1,2p+1}$) are statistically similar. This is inferred by testing the hypothesis $H_0 : R_{m,p} = 0.5$ against the alternative $H_a : R_{m,p} \neq 0.5$, with a pre-set type I error α, using Fisher's exact test. If the null hypothesis is not rejected, that split $R_{m,p}$ is set to 0.5 and the adjacent hazard increments are considered equal and the time bins declared "fused". The hypothesis testing can be applied to all M levels of the tree or just a higher resolution subset of the tree. Because bins are fused *a priori*, this method can then reduce the number of parameters sampled in the MCMC routine, possibly decreasing computation time and increasing the robustness of the estimator.

2.3 Estimation in the MRH model

Estimation is performed in two steps: the pruning step and the MCMC steps. The pruning step is run only once for each of the $\mathcal{L}$ hazard rates at the beginning of the algorithm as a pre-processing step in order to finalize the MRH tree priors. The $R_{m,p;\ell}$ parameters for which the null hypothesis is not rejected are set to 0.5 with probability 1, while the rest are estimated in the Markov chain Monte Carlo (MCMC) routine. Details on the MCMC routine and prior values can be found in [4, 7] and [11].

3 Tongue cancer data

To demonstrate the different features of the package, we use the "tongue" data set available in the R data set available in the R survival package. The data set contains 80 subjects with tongue cancer who had a paraffin-embedded sample of the cancerous tissue taken at the time of surgery, with survival times recorded for each patient (in weeks), as well as the tumor DNA profile (aneuploid or diploid) (see [26] for details). The study went for 400 weeks, with a median survival time equal to 69.5 weeks (SD = 67.3), and 33.8% of the subjects were censored. Between 250 and 350 weeks there were zero failures (censored or uncensored), and after 200 weeks there were zero uncensored failures. This data set is also presented in [18] and analyzed [22]. Table 1 summarizes the data in more detail.

In the code shown in Section 4, the tongue data is named "tongue", and has the following variable names:

- time: Survival time (in weeks) from time of surgery.

Table 1: Sample characteristics of 80 patients in the "tongue" data set found in the R survival package, stratified by tumor type.

	Aneuploid		Diploid		Total Sample	
	N	%	N	%	N	%
Uncensored	31	59.6	22	78.6	53	66.25
Censored	21	40.4	6	21.4	27	33.75
Total	52	65.0	28	35.0	80	100.0

- type: Indicator denoting which tumor group the patient belongs to ('1' = aneuploid, '2' = diploid).
- delta: The censoring indicator, which equals '1' if the failure is observed, and '0' if right-censored.

4 Using MRH

In this section, we provide code and discussion on how to analyze survival data using the MRH package, demonstrating analyses using the tongue data to quantify survival times post-surgery.

4.1 Selecting the time resolution (specifying M)

Because the MRH methodology is based on a binary partition, it divides the total study time into $J = 2^M$ time intervals (or "bins"). The first step to fitting an MRH model in R is to determine how many bins are required. (The current version of the MRH package assumes all bins are of equal length.) The choice of M is can be determined through biological rationale or using penalized likelihood criteria (such as DIC, see [28]). In some cases the user might wish to explore how different choices of M affect the bin lengths and implications on the model interpretation. In these instances, the `FindBinWidth()` function can help. Below, the bin width is calculated for values of M ranging from 2 to 10 for different time units (seconds, minutes, hours, days, weeks, months, years). The user must provide the vector of survival times, and specify the original unit of the survival times (seconds ('s'), minutes ('m'), hours ('h'), days ('d'), weeks ('w'), months ('m'), and years ('y')).

```
data(tongue)
```

```
 FindBinWidth(time = tongue$time, delta = tongue$delta, time.unit = 'w')
[1] The mean failure time is  73.825 weeks
[1] The median failure time is  69.5 weeks
[1] The range of failure times is  1 to 400 weeks
```

```
        secs      mins     hours       days      weeks      months       years
M2  60480000 1008000.0 16800.000 700.000000 100.000000 22.99794661 1.916495551
M3  30240000  504000.0  8400.000 350.000000  50.000000 11.49897331 0.958247775
M4  15120000  252000.0  4200.000 175.000000  25.000000  5.74948665 0.479123888
M5   7560000  126000.0  2100.000  87.500000  12.500000  2.87474333 0.239561944
M6   3780000   63000.0  1050.000  43.750000   6.250000  1.43737166 0.119780972
M7   1890000   31500.0   525.000  21.875000   3.125000  0.71868583 0.059890486
M8    945000   15750.0   262.500  10.937500   1.562500  0.35934292 0.029945243
M9    472500    7875.0   131.250   5.468750   0.781250  0.17967146 0.014972621
M10   236250    3937.5    65.625   2.734375   0.390625  0.08983573 0.007486311
```

While there are many acceptable bin widths for an analysis, we generally aim to use one that will allow us to observe a maximum amount of detail in the hazard rate while still remaining computationally feasible and biologically plausible. In a typical analysis, we would be inclined to use a model with $M = 6$, creating bins that are 6.25 weeks long. However, for the purposes of this manuscript, we reduce the number of bins to 16 (i.e. $M = 4$) to reduce computing time while still allowing for adequate demonstration of the package.

4.2 Fitting the MRH model

Once M has been determined, the MRH model is fit using `estimateMRH()`. In the examples below, M is equal to 4, with each bin representing 25 weeks. We examine MRH models including the treatment covariate under the proportional and non-proportional hazards assumptions, as well as with various levels of pruning.

4.2.1 Proportional hazards model

In this example, we include the tumor-type in the model under the proportional hazards assumption. The code to fit the model is below:

```
fit.PH = estimateMRH(Surv(time, delta) ~ type, data = tongue,
M = 4, maxStudyTime = 400, outfolder = 'MRHresults_PH')
```

The user is notified of the approximate running time and after every 5,000 iterations completed by the routine. The output for each model is placed into a sub-folder created by the `estimateMRH()` routine within the working directory. The default folder is named "MRHresults", however, in this example we have specified the output folder as "MRHresults_PH" through the "outfolder" option. (The folder name can also be added to a pathname, but the path must be accessible from the working directory.) Because parameters are estimated through MCMC sampling, convergence of the fitted model should be assessed after the routine has finished (see Section 4.6 for details).

The fitted model can be examined using both the `summary.MRH()` function as well as the `plot.MRH()` function (see Figure 1 for plot results):

```
# Get the estimated model results
results = summary(fit.PH)
names(results)
[1] "hazardRate"       "beta"             "SurvivalCurve"    "CumulativeHazard"
[5] "d"                "H"                "Rmp"
results$hazardRate
            hrEst    hrq.025    hrq.975
h.bin1  0.00657792 0.00261944 0.01567524
h.bin2  0.00484148 0.00178080 0.01202800
h.bin3  0.00457100 0.00162196 0.01197604
h.bin4  0.00300732 0.00077372 0.00966172
h.bin5  0.00709884 0.00195920 0.02141968
...
# Plot the results (default, smoothed version, cumulative hazard, survival curve)
plot(fit.PH)
plot(fit.PH, smooth.graph = TRUE, smooth.df = 10)
plot(fit.PH, plot.type = 'H')
plot(fit.PH, plot.type = 'S')
```

(Note that in this manuscript, we may use '...' when showing output to indicate that further output is available, but not shown for the legibility purposes.)

4.2.2 Non-proportional hazards model

In many studies with long-term follow-up, the effects of certain covariates change over time, violating the proportional hazards assumption. A figure of the Kaplan-Meier curves ([17]) (see Figure 2, left) show evidence that the effects of treatment may not be proportional. In addition, we can examine if the treatment effect seems to be proportional by creating a cox model (using `coxph()` in the survival package) and then examine the Schoenfeld residuals and smoothed line (see [9]), which should be straight if the effects of tumor type were proportional (see Figure 2, right). Results from these figures provide evidence that the treatment effect is not proportional.

Because of this evidence of non-proportionality between the two treatment group hazard rates, we modify the model above using the `nph()` function in the formula portion used to fit the model:

```
fit.NPH = estimateMRH(Surv(time, delta) ~ nph(type), data = tongue,
M = 4, maxStudyTime = 400, outfolder = 'MRHresults_NPH')
```

In creating a model with at least one NPH covariate, it is important to note that the non-proportional covariate must be a nominal categorical variable. More than one non-proportional covariate can be entered in to the model, using repeated `nph()` functions in the formula, and the routine will merge the NPH variables into a single interaction variable, jointly estimating stratified hazard rates for each combination of levels. See `vignette("MRH")` for more information.

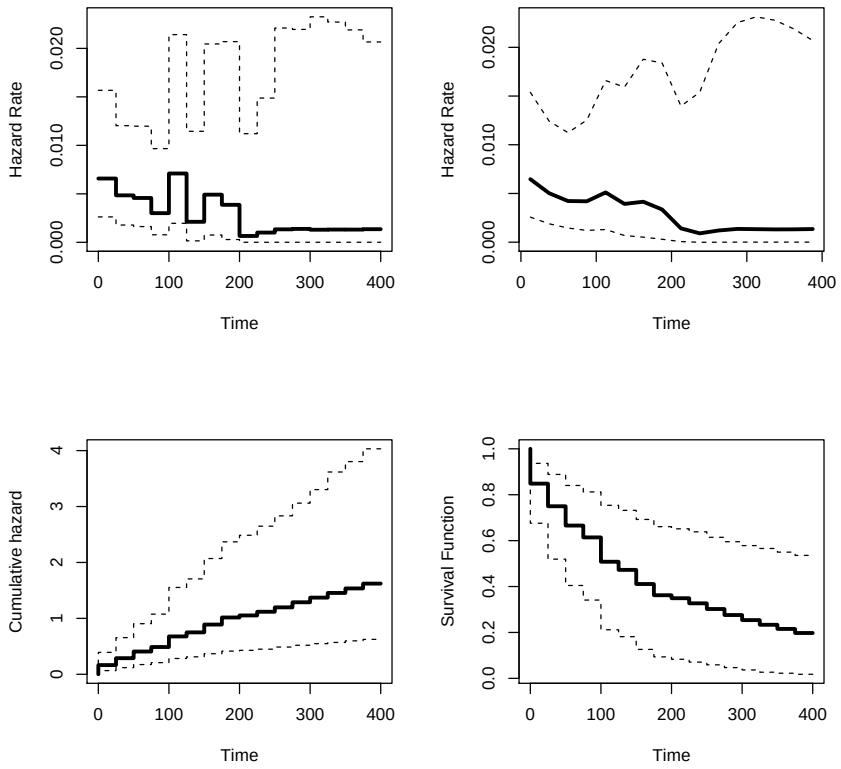

Figure 1: Four graphs of the fitted PH model "fit.PH", all created using the `plot.MRH()` function and associated options on the MRH fitted object. TOP LEFT: Default graph of the hazard rate of death from tongue cancer (`plot(fit.PH)`). TOP RIGHT: Smoothed graph of the hazard rate of death from tongue cancer (`plot(fit.PH, smooth.graph = TRUE, smooth.df = 10)`). BOTTOM LEFT: The cumulative hazard (`plot(fit.PH, plot.type = 'H')`). BOTTOM RIGHT: The survival function (`plot(fit.PH, plot.type = 'S')`).

Summaries for the NPH fitted model are displayed using `summary.MRH()`, with separate estimates provided for each hazard rate, as well as the estimated log-ratio between the hazard rates:

```
summary(fit.NPH)$hazardRate
                hrEst hrq.025 hrq.975
                   hrEst    hrq.025    hrq.975
h.bin1.group1  0.00824948 0.00415240 0.01457768
h.bin2.group1  0.00727956 0.00315452 0.01440876
h.bin3.group1  0.00694756 0.00281784 0.01407212
h.bin4.group1  0.00718136 0.00223352 0.01665544
h.bin5.group1  0.00716788 0.00138364 0.02265908
```

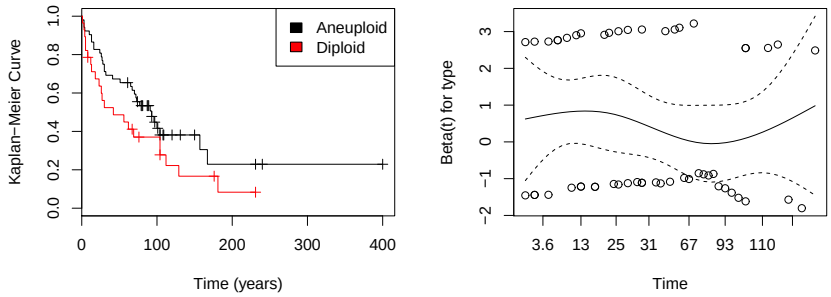

Figure 2: LEFT: Kaplan Meier survival curves of the two treatment groups. The two survival curves do not appear proportional to one another, particularly as they get closer in value towards the end of the study. In addition, each '+' sign indicates a censored observation, so many subjects do not have observed biochemical failure. RIGHT: A plot of the Schoenfeld residuals (based on work done by [9]) shows a smoothed line that does not appear to be straight, indicating non-proportionality between the hazards. The information provided by these two graphics warrants analyses including the treatment covariate under a non-proportional hazards assumption.

```
...
h.bin12.group2 0.00713068 0.00002044 0.16397088
h.bin13.group2 0.00720296 0.00001700 0.16742536
h.bin14.group2 0.00703252 0.00001728 0.16919512
h.bin15.group2 0.00759180 0.00002364 0.18139816
h.bin16.group2 0.00706220 0.00001868 0.17366552

summary(fit.NPH)$beta
                    betaEst betaq.025 betaq.975
beta.type.2.bin1   0.810746 -0.112749  1.701116
beta.type.2.bin2   0.403418 -0.910105  1.572686
beta.type.2.bin3   0.367961 -1.117286  1.688406
beta.type.2.bin4  -2.127125 -6.461301  0.686255
...
```

Plots of the estimated hazard rates and the log-hazard ratio of the NPH covariate can also be created using `plot.MRH()`. Example graphing code is below (see Figure 3 for plot results):

```
# Plot the default graph (the hazard rate of each treatment group)
plot(fit.NPH)
# Plot the log-hazard ratio of the treatment effect
plot(fit.NPH, plot.type = 'r')
# Plot the hazard rates each on a separate graph
plot(fit.NPH, combine.graphs = FALSE)
```

Note that the plotting options shown in Section 4.2.1 can also be used in the NPH plots. Details on convergence checking, output generation, and other specifications on the MCMC chains are found in Section 4.6.

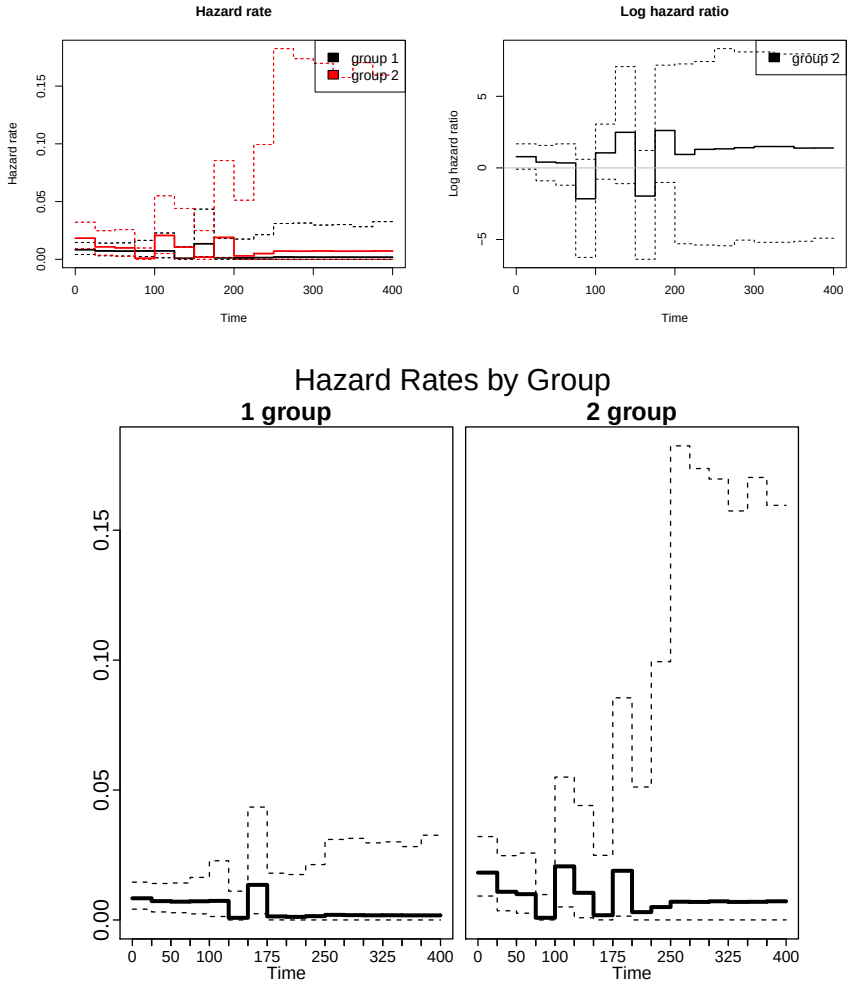

Figure 3: Three graphs of the fitted NPH model, all created using the `plot.MRH()` function and associated options on the MRH fitted object. TOP LEFT: Default graph of the hazard rates of biochemical failure for each treatment group (`plot(fit.NPH)`). TOP RIGHT: Log-hazard ratio of the treatment effect over time (`plot(fit.NPH, plot.type = 'r')`). BOTTOM: The estimated hazard rates for each treatment group, each shown on a separate graph (`plot(fit.NPH, combine.graphs = FALSE)`).

4.3 Pruning the MRH model

As mentioned previously, in instances where the number of observed failures is small, estimates of the hazard rate can be difficult to obtain due to lack of information. In the examination of biochemical failure, the number of subjects who have observed biochemical failure is small, particularly towards the end of the study where a lot of censoring occurs. In this particular data set, only 50% of subjects have an observed biochemical failure, and only 13% of the observed failures occur after 5 years.

In the MRH package, bins can combined using the "pruning" method ([5]), implemented using the "prune" option in the `estimateMRH()` function. It is possible to prune one to all levels of the prior tree based on biological rationale and the degree of possible smoothing the user desires. By default, all levels of the tree are pruned using a significance level $\alpha = 0.05$. However, the number of levels pruned can be adjusted using the "prune.levels" option, and the significance level can be controlled using the "prune.alpha" option. (Note that smaller values of α will lead to smoother prior trees, as the null hypothesis that the two bins are similar will not be rejected as often.)

Below, we show code for pruning the MRH prior tree for the NPH model, combining bins in the bottom 3 levels and combining bins in all levels. In the NPH model, each hazard rate is pruned separately:

```
# Model with no pruning (same as the model shown in Section 4.2)
fit.NPH = estimateMRH(Surv(time, delta) ~ nph(type), data = tongue,
      M = 4, maxStudyTime = 400, outfolder = 'MRHresults_NPH')

# Model with the bottom level of the tree pruned
fit.NPHprune1 = estimateMRH(Surv(time, delta) ~ nph(type), data = tongue,
      M = 4, maxStudyTime = 400, outfolder = 'MRHresults_NPHprune1',
      prune = TRUE, prune.levels = 1)

# Model with bottom 3 levels of the tree pruned
fit.NPHprune3 = estimateMRH(Surv(time, delta) ~ nph(type), data = tongue,
      M = 4, maxStudyTime = 400, outfolder = 'MRHresults_NPHprune3',
      prune = TRUE, prune.levels = 3)
```

Results of the estimates for the three different models can be seen in Figure 4. As expected, the model with no pruning (left graph) has the highest variability among the three models. The model with the bottom 3 levels pruned shows more variability than the model with all levels pruned, however few bins were merged in the top 3 levels of the tree, so results for the two models (shown in the center and right graphs) are similar.

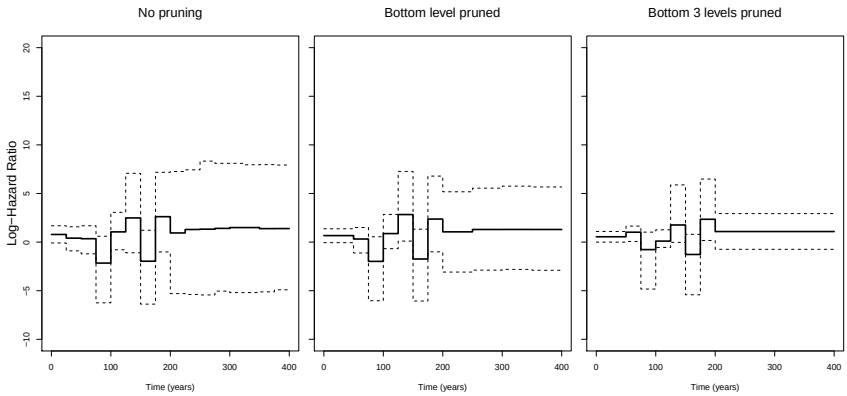

Figure 4: Graph comparing the estimated log-hazard ratio (solid black line) and 95% credible intervals (grey shading) for the three NPH models, each calculated using different levels of pruning. On the left, the model with no pruning shows wider error bounds than the 3-level pruned model (center) and then all-level pruned model (right).

4.4 Prior parameter effects

In the `estimateMRH()` function, the user has the option to adjust the parameters of the prior distribution for the split parameters. The prior for the $R_{m,p}$ parameters is a beta distribution, with the form:

$$R_{m,p} \sim Beta\left(2\gamma_{m,p}k^m a, 2(1-\gamma_{m,p})k^m a\right).$$

It is possible to sample the parameters k and/or γ using the `estimateMRH()` routine. Alternatively, specific fixed values of k and γ can be defined *a-priori* to influence the shape and smoothness of the hazard rate (the default is fixed values equal to 0.5 for all parameters). Details on the role of these hyperparameters can be found in [4, 3, 5, 11]. Below is a brief explanation of the assumptions that stem from using different fixed or sampled values of k and γ.

4.4.1 Comparison and selection of k

In the MRH package, by default k is fixed 0.5, implying zero *a-priori* correlation among the hazard increments. However, k can be fixed with values greater than 0.5, implying the increments are positively correlated *a priori* (typically leading to smoother estimated hazard functions). Alternatively, if k is fixed with a value less than 0.5, the hazard increments are negatively correlated *a priori* and generally cause estimated hazard rates that are less smooth. It is also possible to sample the parameter k if the user desires. The user may change the default value of $k = 0.5$ by entering a value for"k.fixed," with $k \in (0, \infty)$, in the `estimateMRH` function. Below, we show example code for sampling k, and fixing k at 0.1 and 10:

```
# Sample the k parameters for each hazard rate
```

```
fit.NPH.samplek = estimateMRH(Surv(time, delta) ~ nph(type), data = tongue,
      M = 4, maxStudyTime = 400, outfolder = 'MRHresults_NPH_ksample',
      k.fixed = FALSE)

# Fix k at 0.1
fit.NPH.kpt1 =estimateMRH(Surv(time, delta) ~ nph(type), data = tongue,
      M = 4, maxStudyTime = 400, outfolder = 'MRHresults_NPH_kpt1',
      k.fixed = .1)

# Fix k at 10
fit.NPH.k10 = estimateMRH(Surv(time, delta) ~ nph(type), data = tongue,
      M = 4, maxStudyTime = 400, outfolder = 'MRHresults_NPH_k10',
      k.fixed = 10)
```

Note that in the NPH models, a vector of k values can be entered, with one k for each subgroup hazard rate. However, if only one value of k is specified, that value will be used for all hazard rates. Alternatively, if "k.fixed" is set to `FALSE`, the routine will sample the k parameter(s), putting an exponential hyper prior on k. Estimated log-hazard ratios of the treatment covariate from the models coded above can be observed in Figure 5. As anticipated, the estimate from the model where k is sampled (upper right) has the highest variability. In contrast, the model with a very large k value (lower right) has the lowest variability and also a very smooth estimated function.

4.4.2 Comparison and selection of γ_{mp}

The mean of the prior for each split parameter $R_{m,p}$ is $E(R_{m,p}) = \gamma_{m,p}$, with $\gamma_{m,p} \in (0, 1)$. This allows the user to *a priori* "center" the baseline hazard increments in each bin at a desired value (please see [3] for details).

The default value for the γ parameters is 0.5, however the user may change the default values via the "gamma.fixed" option. In general, unless the user has specific information that would help in the adjustment of the γ parameters, we recommend keeping the value fixed at 0.5. In the case of user input, a $\gamma_{m,p}$ value must be specified for each of the split parameters. In NPH models, a matrix of fixed γ values may be entered, with each column representing the desired values for each subgroup hazard rate. However, this is not required; if only a vector is specified, that vector will be used for all hazard rates. Below, we show code for sampling the γ parameters, putting a beta hyper prior on each $\gamma_{m,p}$:

```
fit.NPH.gammasample = estimateMRH(Surv(time, delta) ~ nph(type), data = tongue,
      M = 4, maxStudyTime = 400, outfolder = 'MRHresults_NPH_gsample',
      gamma.fixed = FALSE)

summary(fit.NPH.gammasample)$gamma
             gammampEst gammampq.025 gammampq.975
gammamp1.0_1   0.755960     0.275969     0.972982
gammamp2.0_1   0.513516     0.134032     0.884236
```

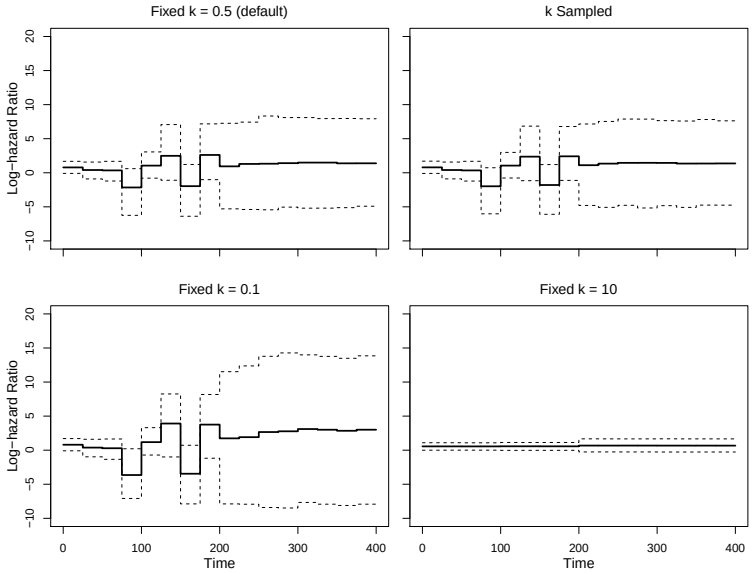

Figure 5: Comparison of the estimated log-hazard ratio of the treatment effect for the NPH model, with k included in the $R_{m,p}$ prior in different ways: fixed at 0.05, implying zero *a priori* correlation between bins (the default value, top left), k sampled (top right), k fixed at 0.1, implying negative correlation between bins (bottom left) and k fixed at 10, implying positive correlation between bins (bottom right). Estimates are shown with solid lines, and 95% credible intervals are shown with dashed lines. As expected, the model using the smallest value of k produces the estimate with the most peaks (i.e. the "bumpiest" estimate), and the model using the largest value of k (equal to 10) produces the smoothes estimate. The model that samples k produces an estimate very close to the default model, with median estimates of k equal to 0.55 for the aneuploid group and 0.56 for the diploid group.

```
gammamp2.1_1   0.461246     0.055678     0.935269
gammamp3.0_1   0.512232     0.113273     0.890278
gammamp3.1_1   0.451540     0.080873     0.877487
...
gammamp4.3_2   0.375792     0.057736     0.854770
gammamp4.4_2   0.478627     0.077879     0.913257
gammamp4.5_2   0.499923     0.077736     0.924315
gammamp4.6_2   0.497371     0.092031     0.917733
gammamp4.7_2   0.502462     0.085781     0.923276
```

The estimated hazard rates and log-hazard ratio can be observed in Figure 6. The estimates do not differ that dramatically, but the variation in the model where the vector of γ parameters is sampled is greater. This is expected, since this model estimates $2^{M-1} \times 2$ more parameters than the default model.

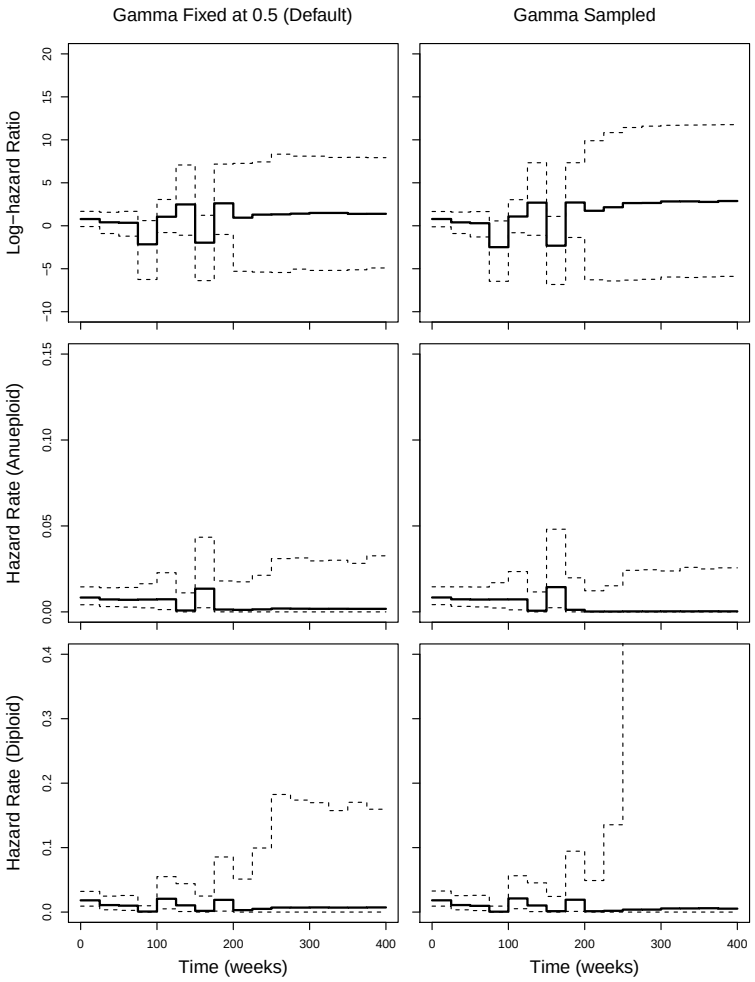

Figure 6: Comparison of the estimated log-hazard ratio and hazard rates (solid lines), and 95% credible intervals (dashed lines) for both treatment groups for the default MRH model (left column) and the MRH modeled with sampled $\vec{\gamma}$ values (right column). While the estimated hazard rates and log-hazard ratios do not differ dramatically, the variation in the model with sampled $\vec{\gamma}$ values is greater. This is expected, as the second model estimates $2^{M-1} \times 2$ more parameters than the default model.

4.5 Model comparison

In comparing different models, it may be desirable to calculate and compare the Deviance Information Criterion (DIC, [28]), the Akaike Information Criterion (AIC, [1]), and the Bayesian Information Criterion (BIC, [25]). In the MRH package, there are two methods for obtaining these values. The fitted model contains the different information criteria values (under the "DIC", "AIC", and "BIC" labels), or the `DIC()` function that can be used on the fitted model or on

Table 2: Comparison of model information criteria values (DIC, AIC, and BIC) for each of the models presented in the previous sections. All values are calculated using the `DIC()` function in the MRH package.

Model			DIC	AIC	BIC
PH	Default		614	676	721
NPH	Default		613	697	783
	Pruned	1 level	608	667	724
		3 levels	604	642	671
	Change k	k sampled	614	701	791
		$k = 0.1$	614	705	790
		$k = 10$	606	690	776
	Sample γ		613	765	922

the chains (see Section 4.7 for details on using the text files of chains). Code for obtaining these values is below, and a comparison of the values for all models shown in this manuscript can be seen in Table 2.

```
fit.NPHprune1$DIC
[1]  607.7004
fit.NPHprune1$AIC
[1] 667.0951
fit.NPHprune1$BIC
[1]  724.2637

DIC(fit.NPHprune1, n = 80)
$neg2loglik.summ
        value
Min.    585.8
1st Qu. 593.2
Median  596.0
Mean    596.5
3rd Qu. 599.3
Max.    619.1

$ICtable
       value
DIC 607.7004
AIC 667.0951
BIC 724.2637
```

4.6 Settings and convergence diagnostics of the MCMC chains

The MRH package allows the user to control different properties of the MCMC chain (i.e. the burn-in, the thinning value, and the maximum number of iterations). In addition, the `estimateMRH()` routine checks for possible convergence,

and can modify the properties of the MCMC chain. Plots are also provided in the output folder that allow the user to assess if the chains have converged. Below, we outline the different methods for fixing the properties of the MCMC chain and assessing convergence.

4.6.1 Setting parameters of the MCMC sampling

The user may use the default values for the burn-in ("burnIn"), thinning value ("thin"), and maximum number of iterations ("maxIter") of the MCMC chains, or they may specify and fix these values. The default maximum number of MCMC iterations is set at 500,000, with a burn-in value of 50,000 and a thinning value of 10. However, based on evidence of chain convergence, these numbers may be changed by the routine unless otherwise specified by the user.

Checking for evidence of convergence

After the first 100,000 MCMC iterations, the chains are checked for autocorrelation and evidence of convergence. (If the maximum number of iterations specified by the user is less than 100,000, then convergence is checked when the maximum number has been reached.) The convergence checking routine is performed via the Geweke diagnostic test ([29]) and the Heidelberger-Welch diagnostic test ([14]) using the `geweke.diag()` and `heidel.diag()` functions available in the coda package ([23]). The convergence algorithm can be seen in the Convergence Algorithm table.

There may be instances in which the user wants to fix `burnIn`, `thin`, or `maxIter` so that the routine does not change these values in the process of checking for convergence, and so that the user is guaranteed the MCMC chain will run `maxIter` times. These values can be fixed (simultaneously or individually) by setting the "fix.burnIn", "fix.thin", or "fix.max" options to TRUE. (By default, these are set to FALSE.)

In addition to the convergence checking performed by the algorithm, diagnostic figures containing trace, density, moving average (calculated by 100), and autocorrelation plots for each parameter are included in the output folder for the model. An example of these plots (for three parameters) can be seen in Figure 7. This graphic allows to user to also assess if convergence has been reached, or if certain parameters are more problematic then others.

Continuing chains

If the routine does not detect evidence of convergence and `maxIter` is reached, then the following message is shown to the user:

```
Warning message:
In estimateMRH(Surv(time, delta) ~ nph(type), data = tongue, M = 4,  :
  Algorithm has not yet converged after  MCMC iterations.
  Parameter estimates may not be reliable.
```

Algorithm Convergence Determine if convergence has been reached and auto-correlation minimized

while The number of iterations is less than `maxIter` **do**

1. Check if thinning value is high enough to reduce autocorrelation: Use `acm()` in the coda package to test the current value of `thin` as well as 5, 10 and 15 times `thin`. The first lag where autocorrelation disappears is the new thinning value. (Note that this value can only be greater than the previous value.)

if New thinning value is different than original value **then**
  Set `thin` to the updated value.
  Thin the stored chains.
end if

2. Convergence check: calculate all p-values for the z-scores returned for each parameter from the `geweke.diag()` function and calculate the results of the Heidelberg-Welch test.

if All p-values < 0.005 and the Heidelberg-Welch test passes **then**
  Convergence reached.
else
  while Convergence is not reached and the number of retained iterations (after thinning) is greater than 1000 **do**
    Burn 20,000 more iterations
    Perform convergence check again.
  end while
end if

if Convergence is reached **then**
  End MCMC sampling routine.
  Update the values of `burnIn` and `thin` (if needed).
else
  Set `thin`, `burnIn`, and stored chains to original values (from before Step #1).
  Perform another 100,000 MCMC iterations. (Or the remaining MCMC iterations needed for `maxIter` to be reached, whichever is smaller.)
end if

end while

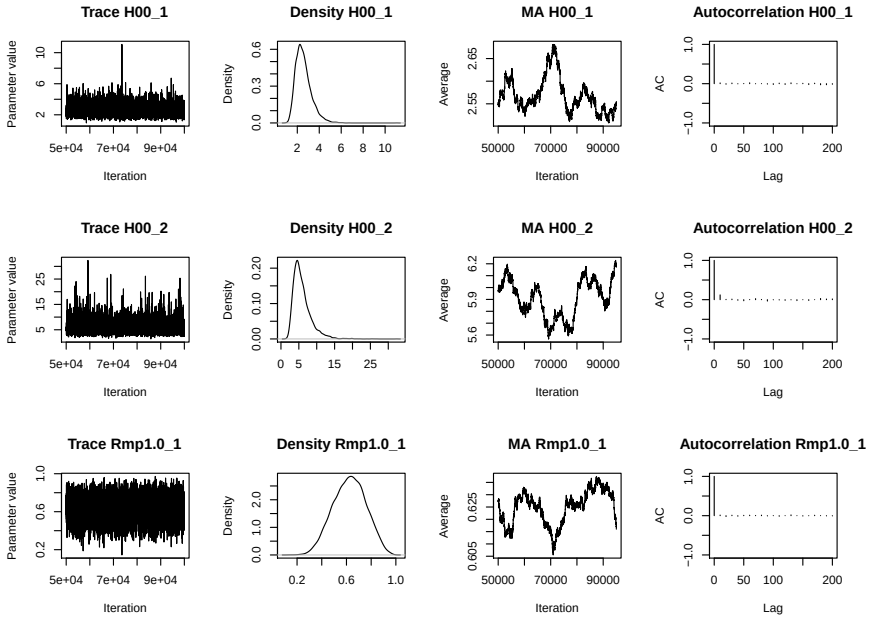

Figure 7: Trace, density, moving average (calculated in increments of 100), and autocorrelation plots for three of the parameters of the NPH 3-level pruned model. (This graph can be found in the "MRHresults_NPH_prune3" output folder, and is titled "convergenceGraphs1.pdf".) This graphic allows to user to also assess if convergence has been reached, or if certain parameters have a more robust estimate than others.

In this instance, the user may want to continue running the MCMC routine using the previously sampled chain values. This can be done using the same call as the in the first model and by setting "continue.chain" equal to TRUE. (Note that this will only work if the output folder name is the same as with the previous model. In addition, if the user specifies new thinning or burn-in values, these will be ignored – the values from the previous chain will be used instead.) If the chain is continued, the routine reads in the chains from the previous set of iterations and initializes the parameters using the last line of retained MCMC values, appending new samples to the previous existing text file store in the output folder. Example code and output for this is below:

```
fit.NPH.continue = estimateMRH(Surv(time, delta) ~ nph(type), data = tongue,
  M = 4, maxStudyTime = 400, outfolder = 'MRHresults_NPH_continue', maxIter = 5000)
....
Warning message:
In estimateMRH(Surv(time, delta) ~ nph(type), data = tongue  :
  Algorithm has not yet converged after  MCMC iterations.
  Parameter estimates may not be reliable.
```

```
fit.NPH.continue = estimateMRH(Surv(time, delta) ~ nph(type), data = tongue,
  M = 4, maxStudyTime = 400, outfolder = 'MRHresults_NPH_continue', maxIter = 50000)

 MCMC routine running.  Calculating estimated runtime for 50000 iterations...

Estimated total run time is 20 minutes

To shorten the run time, re-run with fewer iterations or a smaller number of bins.
```

4.6.2 Gelman-Rubin diagnostic testing

To provide robust estimates of the hazard rate and covariate effects, it is common to run the MCMC routine for the same model multiple times, using the results from all chains to produce estimates and to check for convergence. The MRH package provides a number of options that allow the user to do this more easily.

In the `estimateMRH()` routine, the user can set the "GR" (Gelman-Rubin) option equal to TRUE (by default this is FALSE). In doing this, the routine automatically fixes the chain parameters (`thin`, `burnIn`, and `maxIter`) to what has been specified by the user in the function call, and also adjusts the initial values of the parameters to cover the parameter space. (These are necessary qualifications for the use of the Gelman-Rubin convergence test.)

After all MCMC chains have been sampled, `AnalyzeMultiple()` is available for analyzing multiple MCMC chains for the same model. The `AnalyzeMultiple()` function accepts multiple chains as an input parameter, and then returns to the user the estimated parameter values and α-level credible intervals and the Gelman-Rubin diagnostic information. Within each set of chains, the median, $\alpha/2$%-tile, and $1-\alpha/2$%-tile of the marginal posterior distribution is calculated for each parameter. Then, the median of the medians and percentiles is calculated, and these are the numbers reported as the estimates and credible intervals for each parameter. Example code for this is below:

```
# Fit the three models separately
fit.NPHprune31 = estimateMRH(Surv(time, delta) ~ nph(type), data = tongue,
      M = 4, maxStudyTime = 400, outfolder = 'MRHresults_NPHprune31', GR = TRUE,
prune = TRUE, prune.levels = 3)

fit.NPHprune32 = estimateMRH(Surv(time, delta) ~ nph(type), data = tongue,
      M = 4, maxStudyTime = 400, outfolder = 'MRHresults_NPHprune31', GR = TRUE,
prune = TRUE, prune.levels = 3)

fit.NPHprune33 = estimateMRH(Surv(time, delta) ~ nph(type), data = tongue,
      M = 4, maxStudyTime = 400, outfolder = 'MRHresults_NPHprune31', GR = TRUE,
prune = TRUE, prune.levels = 3)
```

An examination of the initial starting values shows that the initialized values cover the parameter space (which can be accessed with `fit.NPH1$initialValues`).

```
# Get the parameter estimates and credible intervals using all three chains.
# Also check for convergence using the Gelman-Rubin diagnostic test.
results = AnalyzeMultiple(fileNames = c('MRHresults_NPHprune31/MCMCchains.txt',
      'MRHresults_NPHprune32/MCMCchains.txt',
      'MRHresults_NPHprune33/MCMCchains.txt'), maxStudyTime = 400)
names(results)
[1] "hazardRate"       "beta"             "SurvivalCurve"    "CumulativeHazard"
[5] "d"                "H"                "Rmp"              "gelman.rubin"
results$gelman.rubin
          Scale Reduction Factor
H00_1                          1
H00_2                          1
Rmp1.0_1                       1
Rmp4.2_1                       1
Rmp4.3_1                       1
Rmp1.0_2                       1
Rmp4.1_2                       1
Rmp4.3_2                       1
a_1                            1
a_2                            1
lambda_1                       1
lambda_2                       1
```

The results of the Gelman-Rubin diagnostic test can be used to determine if convergence has been reached. If the scale reduction factor is "far from" 1, then it is recommended that more iterations be performed (using the "continue.chain" option).

4.7 Using the MCMC text files

The computation time for convergence of the MRH model may be longer than the user would like to spend waiting (particularly in cases where the number of bins and/or the number of subjects is high). In these instances, the user may want to run the model as a background job, in which case the fitted model results will not be available in the console. Under these circumstances, the user can read in the MCMC chains from the output folder, convert them to an MRH object, and then may use existing functions for plotting and summarizing the chains. The only difference between the usage of the functions is that the maximum study time ("maxStudyTime") must be entered for accurate estimates:

```
# Read in the file from the output folder
mcmc.NPH.prune3 = read.table('MRHresults_NPH_prune3/MCMCchains.txt', header = TRUE)

# Convert to an MRH object
 MRH.NPH.prune3 = as.MRH(mcmc.NPH.prune3)
class(MRH.NPH.prune3)
```

```
[1] "MRH"

#Summarize and plot the results (warning occurs if maximum study time is not entered)
summary(MRH.NPH.prune3)
Error in summary.MRH(MRH.NPH.prune3) :
  Maximum study time (maxStudyTime) needed for hazard rate calculation.
  The maximum study time can be found in the MCMCInfo.txt file in the output folder.
names(summary(MRH.NPH.prune3, maxStudyTime = 400))
[1] "hazardRate"       "beta"             "SurvivalCurve"    "CumulativeHazard"
[5] "d"                "H"                "Rmp"

plot(MRH.NPH.prune3, maxStudyTime = 400)
```

5 Discussion

In this manuscript, we have highlighted the main features of MRH package and demonstrated its use on the tongue cancer data set. Use of the pruning tool and the accommodation of non-proportional hazards makes this package idea for estimation of right-censored survival outcomes when the number of observed failures is small and/or the follow-up period is long.

There are a few other packages that provide a non- or semi-parametric estimate of the hazard rate, and that accommodate non-proportional hazards. Namely, these are `bayesSurv` ([19]), `DPpackage` ([15]), and `timereg` ([24]). While these packages provide their own unique strengths to the analysis of right-censored survival data, the covariate interpretations for all three models are different than that of the MRH model: `bayesSurv` implements AFT survival models, `LDDPsurvival()` in the `DPpackage` package estimates covariates in an ANOVA-like fashion using a Dirichlet Process prior, and the `timecox()` function in the `timereg` package produces cumulative covariate estimates. (For a thorough comparison of these packages, see [13].) The MRH package provides a useful tool for estimation of the hazard rate and covariate effects.

6 Acknowledgements

This work was supported by grants nsf-geo 1211668 and nsf-deb 1316334. The project utilized the Janus supercomputer, which is supported by the National Science Foundation (award number CNS-0821794) and the University of Colorado - Boulder. The Janus supercomputer is a joint effort of the University of Colorado - Boulder, the University of Colorado - Denver, and the National Center for Atmospheric Research. Janus is operated by the University of Colorado - Boulder. The authors thank Yuanting Chen and the researchers at NCAR and the REACCTING project for their helpful advice and input.

References

[1] Hirotugu Akaike. A new look at the statistical model identification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19:716–723, 1974.

[2] P. Andersen, O. Borgan, R. Gill, and N. Keiding. *Statistical Methods Based on Counting Processes*. Springer-Verlag, Berlin, 1993.

[3] P. Bouman, J. Dignam, V. Dukic, and X.L. Meng. A multiresolution hazard model for multi-center survival studies: Application to Tamoxifen treatment in early stage breast cancer. *Journal of the American Statistical Association*, 102:1145–1157, 2007.

[4] P. Bouman, V. Dukic, and X.L. Meng. Bayesian multiresolution hazard model with application to an AIDS reporting delay study. *Statistica Sinica*, 15:325–357, 2005.

[5] Y. Chen, Y. Hagar, J. Dignam, and V. Dukic. Pruned Multiresolution Hazard (PMRH) models for time-to-event data. *Bayesian Analysis*, In Review, 2014.

[6] J. Dignam, V. Dukic, S. Anderson, E. Mamounas, D. Wickerham, and N. Wolmark. Hazard of recurrence and adjuvant treatment effects over time in lymph node-negative breast cancer. *Breast Cancer Research and Treatment*, 116:595–602, 2009.

[7] V. Dukic and J. Dignam. Bayesian hierarchical multiresolution hazard model for the study of time-dependent failure patterns in early stage breast cancer. *Bayesian Analysis*, 2:591–610, 2007.

[8] T. Ferguson. Prior distributions on spaces of probability measures. *The Annals of Statistics*, 2:615–629, 1974.

[9] P. Grambsch and T. Therneau. Proportional hazards tests and diagnostics based on weighted residuals. *Biometrika*, 81:515–526, 1994.

[10] Y. Hagar, D. Albers, R. Pivavarov, H. Chase, V. Dukic, and N. Elhadad. Survival analysis with Electronic Health Record data: Experiments with Chronic Kidney Disease. *Statistical Analysis and Data Mining*, 7:385–403, 2014.

[11] Y. Hagar, J. Dignam, and V. Dukic. Modeling long-term outcomes and treatment effects after androgen deprivation therapy for prostate cancer. *In review*, http://arxiv.org/abs/1509.01275, 2015.

[12] Y. Hagar and V. Dukic. *MRH package in R*, 2014.

[13] Y. Hagar and V. Dukic. Comparison of hazard rate estimation in R. *Submitted*, page http://arxiv.org/abs/1509.03253, 2015.

[14] P. Heidelberger and P.D. Welch. A spectral method for confidence interval generation and run length control in simulation. *Communication of the ACM*, 24:233–245, 1981.

[15] A. Jara, T. Hanson, F. Quintana, P. M′ueller, and G. Rosner. *Package DPpackage*, 2006.

[16] J.D. Kalbfleisch and R.L. Prentice. *The Statistical Analysis of Failure Time Data*. SWiley, Chichester, 2002.

[17] E. Kaplan and P. Meier. Nonparametric estimation from incomplete observations. *Journal of the American Statistical Association*, 53:457–481, 1958.

[18] J.P. Klein and M.L. Moeschberger. *Survival analysis: techniques for censored and truncated data*. Springer, New York, 2003.

[19] A. Komárek. *Package bayesSurv*, 2004.

[20] M. Lavine. Some aspects of Polya tree distributions for statistical modelling. *The Annals of Statistics*, 20:1222–1235, 1992.

[21] P. Müller and A. Rodriguez. *Nonparametric Bayesian Inference*. Institute of Mathematical Statistics and American Statistical Association, Beachwood, Ohio and Alexandria, Virginia, USA, 2013.

[22] P. Müller and A. Rodriguez. *Nonparametric bayesian inference*. Project Euclid, Ohio, USA, 2013.

[23] M. Plummer, N. Best, K. Cowles, and K. Vines. CODA: Convergence diagnosis and output analysis for MCMC. *R News*, 6(1):7–11, 2006.

[24] T. Scheike. *timereg package in R*, 2008.

[25] G. E. Schwarz. Estimating the dimension of a model. *Annals of Statistics*, 6:461–464, 1978.

[26] B.J. Sickle-Santanello, W.B. Farrar, J.F. Decenzo, S. Keyhani-Rofagha, J. Klein, D. Pearl, H. Laufman, and R.V. O'Toole. Technical and statistical improvements for flow cytometric dna analysis of paraffin-embedded tissue. *Cytometry*, 9:594–599, 1988.

[27] D. Sinha and D. K. Dey. Semiparameteric Bayesian analysis of survival data. *Journal of the American Statistical Association*, 92:1195–1212, 1997.

[28] D. J. Spiegelhalter, N. G. Best, B. P. Carlin, and A. Van der Linde. Bayesian measures of model complexity and fit (with discussion). *JRSSB*, 64:583–616, 2002.

[29] J. Geweke (1992) Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments. *Bayesian Statistics 4*, (J.M. Bernardo, J.O. Berger, A.P. Dawid and A.F.M. Smith, eds) 169-193. Oxford Univ. Press.